CLAIMS

WHAT IS CLAIMED IS:

1. A communications system, comprising:

a first queue pair (QP) associated with a first connection, the first QP comprising a first send queue (SQ);

a second QP associated with a second connection, the second QP comprising a second SQ; and

a general pool comprising a shared receive queue (SRQ), the SRQ being shared by the first QP and the second QP.

2. The communications system according to claim 1, wherein the first QP and the second QP are part of a particular node.

3. The communications system according to claim 1, wherein the SRQ comprises an amount of resources that is statistically determined.

4. The communications system according to claim 3, wherein the amount of resources is statistically determined based upon empirical resource requirements of the first connection and the second connection.

5. The communications system according to claim 3, wherein the amount of resources is statistically determined based upon dynamic statistics.

6. The communications system according to claim 1, wherein the SRQ comprises a plurality of posted buffers.

7. The communications system according to claim 1, wherein the first QP does not have its own dedicated received queue (RQ).

8. The communications system according to claim 1, wherein the first QP conducts communications over the first connection as if the first QP has access to more resources of the SRQ than a statistical determination of resource requirements of the first connection.

9. The communications system according to claim 1,

wherein the first QP comprises a first limit queue (LQ) that limits an amount of resources of the SRQ that the first QP can access, and

wherein the second QP comprises a second limit queue (LQ) that limits an amount of resources of the SRQ that the second QP can access.

10. The communications system according to claim 9, wherein if the first QP exceeds a limit as set forth in the first LQ, then the first connection is dropped.

11. The communications system according to claim 9, wherein the first LQ and the second LQ are managed locally.

12. The communications system according to claim 11, wherein the first LQ and the second LQ are managed locally without communications with other nodes.

13. The communications system according to claim 9, wherein at least one of the first LQ and the second LQ is a soft limit.

14. The communications system according to claim 13,

wherein, if the soft limit is reached, then a connection behavior of the first connection or the second connection is analyzed before a response is generated, and

wherein the generated response is based on the analyzed connection behavior.

15. The communications system according to claim 9, wherein at least one of the first LQ and the second LQ is a hard limit.

16. The communications system according to claim 15, wherein, if the hard limit is reached, then an automatic response is generated.

17. The communications system according to claim 15, wherein at least one of the first LQ and the second LQ comprises a maximum value relating to a size of a hole in a transport protocol sequence space.

18. The communications system according to claim 9, wherein the limit relates to at least one of an in-order message, an out-of-order message, a segment of the in-order message and a segment of the out-of-order message.

19. The communications system according to claim 1, wherein resource allocation for the SRQ is managed locally.

20. The communications system according to claim 1,
wherein the first QP is associated with a first completion queue (CQ), and
wherein the second QP is associated with a second CQ.

21. The communications system according to claim 1, wherein the general pool comprises a shared CQ (SCQ), the SCQ being associated with the first QP and the second QP.

22. The communications system according to claim 1, wherein the general pool comprises a memory translation and protection table (TPT) associated with resources of the general pool.

23. The communications system according to claim 1, wherein at least one of the first connection and the second connection is a remote direct memory access (RDMA) connection.

24. The communications system according to claim 1, wherein the at least one of the first connection and the second connection is an Internet small computer system interface (iSCSI) over RDMA (iSER) connection.

25. A communications system, comprising:

a network interface card interface (NI) comprising a network interface card (NIC) and a NIC driver, the NIC being coupled to the NIC driver; and

a consumer coupled to the NI,

wherein the NIC comprises a first queue pair (QP), a second QP and a shared receive queue (SRQ), and

wherein the first QP and the second QP share the SRQ.

26. The communications system according to claim 25, the consumer communicates with the NI via verbs.

27. The communications system according to claim 25, wherein the consumer comprises a user space application.

28. The communications system according to claim 25, wherein the consumer comprises a kernel space application.

29. The communications system according to claim 25, wherein at least one of the consumer and the NI comprises an SRQ manager that employs statistical provisioning in managing resources of the SRQ.

30. The communications system according to claim 25, wherein the NIC comprises a remote-direct-memory-access-enabled NIC.

31. The communications system according to claim 25, wherein at least one of the consumer and the NI locally manages provisioning of the SRQ using statistical information.

32. The communications system according to claim 25, wherein the first QP comprises a send queue (SQ) and a limit queue (LQ).

33. The communications system according to claim 25, wherein the consumer comprises a verb consumer.

34. A method for communicating, comprising:

establishing a first connection associated with a first queue pair (QP);

establishing a second connection associated with a second QP;

concurrently sharing a single receive queue (RQ) between the first QP and the second QP; and

provisioning the single RQ using statistical information.

35. The method according to claim 34, comprising:

managing locally resources of the single RQ.

36. The method according to claim 34, comprising:

managing dynamically resources of the single RQ.

37. The method according to claim 34, comprising:

managing locally RQ resources available to the first QP.

38. The method according to claim 37, wherein managing comprises limiting RQ resources available to the first QP.

39. The method according to claim 34, wherein the first QP is used by a user space application.

40. The method according to claim 34, wherein the first QP is used by a kernel space application.